

AD-A069 862

UNIVERSITY COLL LONDON (ENGLAND) DEPT OF STATISTICS --ETC F/6 17/2
UNIVERSITY COLLEGE LONDON INDRA PROJECT.(U)
MAR 79 P T KIRSTEIN

N00014-77-6-0005

UNCLASSIFIED

TR-55


NL

| OF |
AD
A069862



END
DATE
FILMED

7-79
DDC



NATIONAL BUREAU OF STANDARDS
MICROCOPY RESOLUTION TEST CHART

LEVEL

INDRA TECHNICAL
REPORT TR-55

(14)

INDRA
TECHNICAL
REPORT

A044419
A025-203

A069862

(6) UNIVERSITY COLLEGE LONDON
INDRA PROJECT

DD
REF ID:
JUN 18
RECEIVED
KFC

(9) QUARTERLY STATUS REPORT, NO. 17,
1 JULY - 30 SEP 1978

by

(10) Peter T. Kirstein

(11) March 1979

This document has been
for public release and no
distribution is unlimited.

(12) 8p.

DOC FILE COPY

Submitted to:

Defence Advanced Research Projects Agency,
1400 Wilson Boulevard
Arlington,
Virginia 22209.

Attn. Dr. Vinton G. Cerf

This research was partially supported by the US Office
of Naval Research under Contract N00014-77-G-0005

(15)

Dept of Statistics & Computer Science
University College, London

409 714 Lur 79 05 17

1. INTRODUCTION

This report serves as a quarterly starting report on the following project: British Library (SI/G/ 172 and 221), Ministry of Defence (2047/77), Science Research Council (GR/A/54168) and US Office of Naval Research (N00014-77-G-0005).

→ This quarter saw the first major change on our ARPANET-EPSS-RL convention in several years, with the introduction of a completely new release of the PDP-9 operating system, (Section 2). X.25 activities made slow progress on the experimental front (Section 3), but a number of significant paper studies were made. The SATNET activity came to a very rewarding finale, (Section 4). The SATNET Experiment officially terminated at the end of the quarter. We were able to make a number of interesting User Level Experiments, using the tools which had been developed in the previous couple of years.

Most of Internetwork activity during the quarter was again in paper studies, (Section 5). Of particular importance were the study of how best to provide Transnet File Transfer Facilities, and a study of the X.25 CCITT protocol. Our ~~One~~ ^{The} simulation work (Section 6) produced two publications. ~~One~~ ^{The} ^{first} was a study of Hop-by-Hop versus end-to-end acknowledgement; the second contrasted the case of positive and negative acknowledgement. Finally, in our facsimile work, we learnt how to control our DACOM machine from the same hardware device as is used to drive an HDLC interface.

Accession For	✓
NTIS G-141	
DEC TAB	
Unannounced	
Justification	
By	<i>on files</i>
Distribution	
Availability Codes	
Availand/or	
special	
ist	<i>A</i>

79 05 17 023

2. MULTIMACHINE SYSTEM

The version of the operating system under SWITCH, which included EPSS-ARPANET protocol conversion for interactive traffic and ARPANET-RL conversion for both terminal and file traffic, was introduced into service early in the quarter. The usual number of problems were encountered at first, particularly on file traffic. Many of those have been cleared, but some remain. The system has now grown so large that it is almost impossible to put together without the 40M byte disk. This is unfortunate, because the disk has not been very reliable. Some problems in character terminal mappings have been uncovered (1), and steps are being taken to remedy them. The RSRE-ARPANET software has not, and will not, be included. RSRE has decided it no longer needs to attach its GEC 4080, and wished to use the ARPANET connection for another project.

3. X.25 ACTIVITIES

The level 2 of X.25 has been recoded in Assembler. The Assembler versions, together with a "Universal" operating system interface, should make it easier to transfer the X.25 code to run under other operating systems. MOS and UNIX are two candidates which we will be studying in due course. The interface specification and frame level have been documented (2), (3).

The INTEL 8080 and PDP 9 versions of the X.25 software have not progressed. The first is being abandoned, because the software is too big to co-exist with the Facsimile software in one 8080. Our hardware does not allow multiprocessor working, and as a group we are developing support facilities mainly for LSI-11s at present.

The second has not progressed because the person working on it has left; he hopes to complete it part-time, but this is likely to be a slow process.

Based on our experience with the LSI-11 implementation of X.25, we have studied the problem of connecting Hosts to X.25 by a "Black Box". On one side it would support X.25 and X.28, on the other it would attach to character terminals ports. While this would be a poor long term solution, it presents an easy way for Hosts to attach X.25 nets with no operating system changes. Some equipment already exists for this purpose. We have analysed its flexibility and suitability (4). No decision has been made to build such a device.

We have studied various aspects of the PTT proposals on X.25 and its related activities. Some of those concerned with the interconnection of networks in the proposed X.75 standards are discussed in (5). Two tutorial papers (6) and (7) were prepared for the NATO Advanced Study Institute on Interlinking

4. SATNET ACTIVITY

During this quarter we began a full set of experiments on SATNET. A new data generation algorithm was also added to the package to permit the simulation and measurement of a high level protocol.

The new algorithm was used to simulate "windowing" bulk data transfers and was loosely based upon TCP. The window size and various retransmission timeouts could be set before the generator began. Extra statistics options were also included to enable the detailed operation of the algorithm to be recorded.

This was a move away from the earlier experiments that used just simple streams of traffic, and a move towards simulating full host-to-host protocols. These experiments were extremely rewarding as they demonstrated not only problems that may occur when this protocol is implemented in high delay networks, but also gave a better insight into the operation of the protocol itself.

The experiment plan was designed to measure the performance of the network for a variety of protocols. We soon discovered that the throughput of the network was severely limited: the maximum throughput was approximately seven packets per second. The measurement plan was quickly modified to determine the reason for this poor performance. Not surprisingly we found that the operating system on the traffic generation machine was causing the bottle-neck. There was very little we could do to increase this throughput. One particular set of experiments surprisingly indicated that the traffic sent from the source seldom arrived in the correct sequence at the destination, even when the packets were sent at relatively infrequent intervals. This was soon isolated to being a software bug in the output queue of the SIMPs, and was subsequently corrected by the Network Control Centre.

The bug in the SIMP code caused the repetition of some of our earlier experiments, as those results were invalidated by the change in SIMP software. The measurement plan was also affected by the low throughput of the ELF operating system. Nevertheless, even with low throughput, we were able to explore delays experienced in SATNET, mainly using the CPODA access algorithm. Two extra timestamp stations located within the SIMPs enabled us to measure delays within the various SIMP routines for the first time.

5. INTERNETWORK ACTIVITIES

The thinking on the development of Protocol Convertors and Port Expanders has continued. The ARPA/INTERNET Programme as a whole is still changing its view on how the gateways should

be configured. Thus a progress report on the configuration thinking at the end of the quarter would, by now, have only historical value.

Two LSI-11s have been configured, one as a Port Expander and one as a TIU Terminal. Both these machines have had the SRI software mounted on them. We had some trouble in making either operational, but succeeded eventually - within the limits of these releases.

The provision of Transnet services has had considerable attention during the quarter. A fairly general study of the types of service which must be provided, and of the tradeoffs necessary in Gateways, is given in (9). A major study of how Bulk Data Transfer could be supported across concatenated computer networks is given in (10). The paper strongly emphasises that it is easier to optimise performance for this application when an end-to-end virtual call service exists - with virtual calls maintained between pairs of Gateways.

Two preparations for the provision of a Transnet File service have been investigated at UCL. The first requiring Protocol translation at a Gateway is our early approach. An ARPANET-EPSS implementation is under development, and made some progress during the quarter; it does not yet work. That implementation runs on the PDP-9 ARPANET-EPSS Gateway. We also expect to complete the implementation, but it will be phased out with the PDP-9s in 1980. The other approach is to provide an end-to-end Network Independent File Transfer (NIFTP) in both EPSS and ARPANET Hosts; Complete Files will then be staged in a CONUS ARPANET PDP-10. Over UCL and its European Nets, the NIFTP will be used; over ARPANET, etc. ordinary

ARPANET FTP could be applied. The implementation of this NIFTP on a TOPS-20, mentioned in the previous quarterly report, has continued.

6. SIMULATION

The group has had a simulation activity for some time, but it is only recently that there has been a publishable output. During this quarter, two papers were prepared. In one (12) the tradeoffs between the Hop-by-Hop and end-point approaches to interconnection networks were studied by simulation techniques. We conclude that an "End-Point Architecture" to connecting networks is likely to demand more resources from Hosts. A Hop-by-Hop approach demands more from Gateways, and could provide inferior service. Another paper (13) contrasted the use of positive and negative acknowledgement schemes. We concluded that the positive acknowledgement is more suitable for end-to-end connections with moderate or high reliability. For considerable connections, negative acknowledgement may be preferable, because this greatly reduces the extent of redundant transmissions - which can be large with just positive acknowledgement.

7. FACSIMILE ACTIVITIES

The main activity during the quarter has been in learning to drive the new DACOM device from an LSI-11. We have found that it is possible to use the same interface as we use for HDLC - but with the protocol run in another Binary Synchronous mode.

REFERENCES

1. Higginson, P.L.: The need for mapping on SWITCH/EPSS terminal terminal services, INDRA 684, 1978.
2. Bradbury, C.: Interface Specification for the X.25 Level 2 (HDLC) modules used on the LSI-11, INDRA 682, 1978.
3. Bradbury, C.: X.25 Frame Level Implementation, INDRA 683, 1978.
4. Higginson, P.L. and P.T. Kirstein : The UCL Black Box for the attachment at BLAISE to X.25 Networks.
5. Grossman G.R. and A.J. Hinchley : Issues in International Public Data Networks, to be published in Proc. USA-Japan Conference, 1, 1978.
6. Hinchley, A.J.: The service aspects of the X.25 interface, to be published in Advanced Study Institute on Interlinking Computer Networks (1979) and INDRA 687, 1978.
7. Kirstein, P.T.: Some international developments in data services, *ibid*, INDRA 685, 1978.
8. Bennett, C.J.: Type of service in the Catenet, INDRA 680. 1978.
9. Bennett, C.J.: Supporting transnet bulk data transfer. Proc. Conf. Flow Control in Computer Networks, eds. Z.I. Grangé and M. Glen, 383, North Holland, 1979.
10. EPSS Higher Level Protocols Group : A Network Independent File Transfer Protocol, INWG Protocol Note 86, 1977.
11. Edge, S.W.: Comparison of Hop-by-Hop and Endpoint approaches to Network Interconnection, Proc. Conf. Flow Control in Computer Networks, eds. J.L. Grangé and M. Glen, North Holland, 359, 1979.
12. Edge, S.W. and A.J. Hinchley : A survey of end-to-end retransmission techniques, *Comp. Comm. Rev.*, 8, 4, 1, 1978.